REMARKS

The Office Action dated May 9, 2006, has been received and carefully considered. In this response, claims 14, 21, and 23-25 have been amended. Entry of the amendments to claims 14, 21, and 23-25 is respectfully requested. Reconsideration of the outstanding rejections in the present application is also respectfully requested based on the following remarks.

I. THE INDEFINITENESS REJECTION OF CLAIMS 24 AND 25

On page 2 of the Office Action, claims 24 and 25 were rejected under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the invention. Applicants have amended claims 24 and 25 to address the Examiner's concerns.

In view of the foregoing, it is respectfully requested that the aforementioned indefiniteness rejection of claims 24 and 25 be withdrawn.

II. THE ANTICIPATION REJECTION OF CLAIMS 1-10, 13-15, AND 23-25

On page 3 of the Office Action, claims 1-10, 13-15, and 23-25 were rejected under 35 U.S.C. § 102(e) as being anticipated by Green et al. (U.S. Patent Application Pub. No. 2003/0167380, hereinafter "Green"). This rejection is hereby respectfully traversed.

Regarding claim 1, the Examiner asserts that Green teaches a storage management system that "copies, prior to execution of each write command, old data present at the at least one unit of storage into the at least one data store" (emphasis added).

Applicants respectfully submit that Green fails to disclose or suggest the copying of old data prior to execution of each write command.

Once implemented, a storage management system according to the claimed invention will perform copy-on-write (COW) operations in response to each and every write command, so that the old data from these individual COW operations can enable a data restoration or recovery to any point in time since the implementation of the storage management system.

Green, however, can only restore data to limited points in time when snapshots were taken. Green lacks the capability to restore data to any point in time because it does not record every change made to the storage system. It is worth noticing that Green relies on pre-scheduled snapshots to back up data for the storage system. *See*, Figures 14-16 and paragraph [0122]. Alternatively, "a snapshot is periodically 'taken' so that a computer system can be restored in the event of failure." Paragraph [0049] (emphasis added). Since Green does not capture all the write commands directed to the storage system, its data

restoration capability is limited. When multiple write commands are directed to the same storage unit within a short period of time between two snapshots, Green's system cannot capture all of those write commands. Indeed, in the time-point-by-time-point analysis of Figure 3 (paragraphs [0064]-[0073]), Green acknowledges that data "H" between the second and third snapshots is lost. "Because address 4 of the volume changed twice between snapshots, only the starting and ending value of this address are captured by the snapshots." Paragraph [0069].

Green also fails to teach or suggest "wherein a record of the old data is timestamped" as recited in claim 1. Green does not teach timestamping each record of the old data overwritten in response to a write command. It is true that Green speaks of "copy on write." Paragraph [0067]. Despite the use of the same term, the so-called "copy on write" disclosed in Green is quite different from the copy on write described in the present application.

First, Green's "copy on write" is conditional and not performed in response to each and every write command. Rather, the old data about to be overwritten is copied if, and only if, it was in the volume when a previous snapshot was taken. *See*, paragraph [0069].

Second, Green's "copy on write" operations do not record timestamps, and, even if they did, such timestamps would be discarded as soon as a snapshot cache is "permanently fixed." The timeline in Figure 3 is provided to facilitate a time-point-by-time-point illustration. Green does not teach that each of the data copied to the snapshot caches has its own timestamp. The snapshot specific cache grids 342, 344 and 346 are just one way of illustrating the changes of content in those snapshot caches. *See*, paragraph [0061]. The real content of each snapshot cache is shown in table 360 in the lower left corner of Figure 3. *See also*, paragraphs [0062]-[0063]. In table 360, each granule value has a corresponding "Address" and "Snapshot No." Because Green cannot and does not even attempt to restore data to any point in time other than the snapshot time points, it is unnecessary for Green to provide a timestamp for old data backed up in each "copy on write."

It should have become clear now that Green uses "copy on write" only as an intermediate step in preparation for each snapshot. Such "copy on write" operations are conditional, have no individual timestamps, and are collapsed into the snapshots as soon as the snapshot caches are permanently fixed.

It is believed that the distinctions described above set claim 1 apart from Green. Green cannot anticipate claim 1

because Green fails to teach or suggest all the elements recited in claim 1.

Regarding claim 14, the Examiner asserts that the method claimed therein is similarly disclosed by Green as discussed above in reference to claim 1. Applicants have amended claim 14 to clarify that old data are copied "prior to execution of each write command." It is believed that claim 14, as amended, is patentable over Green for at least the same reasons why claim 1 is patentable.

Regarding claim 23, the Examiner asserts that the computer readable medium claimed therein is similarly disclosed by Green as discussed above in reference to claim 1. Applicants have amended claim 23 to clarify that old data are copied "prior to execution of each write command." It is believed that claim 23, as amended, is patentable over Green for at least the same reasons why claims 1 and 14 are patentable.

Since claims 2-13 and 26-27 all depend from claim 1, claims 15-20 and 28 all depend from claim 14, and claims 24-25 and 29 all depend from claim 23, these dependent claims should be patentable over Green for at least the same reasons.

In view of the foregoing, it is respectfully requested that the aforementioned anticipation rejection of claims 1-10, 13-15, and 23-25 be withdrawn.

III. THE OBVIOUSNESS REJECTION OF CLAIMS 11, 12, 21, 22, 16-20, AND 26-29

On page 7 of the Office Action, claims 11, 12 and 16-20 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Green in view of "UNIX In A Nutshell" by Daniel Gilly and the staff of O'Reilly & Associates, Inc. (hereinafter "Gilly"). On page 9 of the Office Action, claims 21, 22 and 26-29 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Green in view of Applicant's Admitted Prior Art (hereinafter "AAPA").

It is believed that these obviousness rejections have become moot in view of the deficiencies of the primary reference Green as discussed above. Since none of the additional references teach or suggest the claim elements that Green fails to disclose, the combination of Green with these references cannot render any of the pending claims obvious.

In view of the foregoing, it is respectfully requested that the aforementioned obviousness rejection of claims 11, 12 and 16-20 be withdrawn.

IV. CONCLUSION

In view of the foregoing, it is respectfully submitted that the present application is in condition for allowance, and an early indication of the same is courteously solicited. The

Examiner is respectfully requested to contact the undersigned by telephone at the below listed telephone number, in order to expedite resolution of any issues and to expedite passage of the present application to issue, if any comments, questions, or suggestions arise in connection with the present application.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-0206, and please credit any excess fees to the same deposit account.

Respectfully submitted,

Hunton & Williams LLP

By: _____
Thomas E. Anderson
Registration No. 37,063

TEA/ccl:mia

Hunton & Williams LLP
1900 K Street, N.W.
Washington, D.C. 20006-1109
Telephone: (202) 955-1500
Facsimile: (202) 778-2201

Date: July 10, 2006

Appendix A

1. (Previously Presented) A storage management system for backing up digital content of a storage system comprising a plurality of units of storage, wherein the storage management system comprises:

at least one data store;

wherein the storage management system automatically intercepts write commands issued to the plurality of units of storage, each write command comprising an instruction to overwrite at least one unit of storage with new data; and

wherein the storage management system copies, prior to execution of each write command, old data present at the at least one unit of storage into the at least one data store, wherein a record of the old data is timestamped.

2. (Previously Presented) The storage management system of claim 1, wherein the storage system further comprises one or more physical storage devices on which the digital content of the storage system is stored.

3. (Previously Presented) The storage management system of claim 2, wherein an address for accessing the storage system comprises a device identifier and a location identifier.

4. (Previously Presented) The storage management system of claim 3, wherein the device identifier identifies a physical storage device.

5. (Previously Presented) The storage management system of claim 3, wherein the device identifier identifies a logical device.

6. (Previously Presented) The storage management system of claim 1, wherein the digital content of the storage system can be accessed by specifying an address and a time, and wherein the time specifies that the digital data retrieved from the address is the most recent digital data that was written to the address at or before the time.

7. (Previously Presented) The storage management system of claim 6, wherein the time is explicitly specified in a request to access a unit of storage.

8. (Previously Presented) The storage management system of claim 6, wherein the time is specified in a command to the storage system separate from a request to read a unit of storage.

9. (Previously Presented) The storage management system of claim 6, wherein the storage management system creates a virtual device, wherein the time is specified when the virtual device is created, and is applied when the virtual device is accessed.

10. (Previously Presented) The storage management system of claim 9, wherein new data is written to the virtual device without overwriting data that was written to the storage system after the time specified when the virtual device was created.

11. (Previously Presented) The storage management system of claim 6, wherein a command to the storage system specifies that the time is implicitly a current time.

12. (Previously Presented) The storage management system of claim 6, wherein the time is specified relative to a current time.

13. (Previously Presented) The storage management system of claim 1, wherein the units of storage are blocks.

14. (Currently Amended) A method for backing up digital content of a storage system having a plurality of units of storage, the

method comprising:

intercepting, automatically, write commands issued to the storage system, wherein each write command comprises an instruction to overwrite at least one unit of storage with new data; and

copying, prior to execution of ~~the~~ each write command, old data present at the at least one unit of storage into a data store, wherein a record of the old data is timestamped;

whereby digital content of the storage system can be accessed by specifying an address and a time to access the most recent data stored on the storage device at the address at or before the time.

15. (Original) The method of claim 14, wherein the address comprises a device identifier and a location identifier.

16. (Original) The method of claim 14, wherein specifying the time comprises implicitly specifying the time.

17. (Previously Presented) The method of claim 16, wherein implicitly specifying the time comprises sending a command to the storage system to use a current time as the time.

18. (Original) The method of claim 14, further comprising presenting a virtual storage device for which the time is implicitly set to the specified time for all addresses of the virtual storage device.

19. (Original) The method of claim 18, further comprising writing data to the virtual storage device.

20. (Previously Presented) The method of claim 14, wherein specifying the time comprises specifying the time relative to a current time.

21. (Currently Amended) ~~Apparatus~~ An apparatus for storing data, the apparatus comprising:

a storage appliance that interfaces with a computer;

one or more physical storage devices that interface with the storage appliance, the one or more physical storage devices having a plurality of storage units, each such physical storage device controlled by the storage appliance;

wherein the storage appliance comprises at least one current store and at least one time store, the at least one current store maintaining a current mirror copy of digital content in the one or more physical storage devices, and

wherein, each time immediately before a storage unit is overwritten with new data, any old data present at that storage unit is timestamped and stored in the at least one time store;

wherein the storage appliance presents one or more virtual storage devices to the computer based on the at least one current store and the at least one time store, and wherein digital data on each of the virtual storage devices is accessed by specifying an address and a time.

22. (Previously Presented) The apparatus of claim 21, wherein the time specifies that the digital data retrieved from the address is the most recent data that was written to the address at or before the time.

23. (Currently Amended) A computer readable medium having code for causing a processor to control a storage system, the storage system comprising a plurality of units of storage, the computer readable medium comprising:

code adapted to automatically intercept write commands issued to the storage system, wherein each write command comprises an instruction to overwrite at least one unit of storage with new data; and

code adapted to copy, prior to execution of ~~the~~ each write

command, old data present at the at least one unit of storage into a data store, wherein a record of the old data is timestamped;

wherein digital content of the storage system is accessible with a storage device command, the storage device command comprising a storage device address identifying the location of one or more units of storage and a time specification specifying data most recently stored at the storage device address at or before a specified time.

24. (Currently Amended) The computer readable medium of claim 23, wherein the storage device command is a write command and the ~~point in~~ specified time is a current time.

25. (Currently Amended) The computer readable medium of claim 23, wherein the storage device command is a read command and the ~~point in~~ specified time is a past time.

26. (Previously Presented) The storage management system of claim 1, wherein the at least one data store comprises a first data store and a second data store, and wherein the first data store maintains a current mirror copy of digital data stored in the plurality of units of storage, and wherein the second data

store contains the old data and the timestamped record of the old data.

27. (Previously Presented) The storage management system of claim 26, wherein, after the old data is copied to the second data store, the at least one unit of storage is overwritten with the new data, and the current mirror copy in the first data store is updated with the new data.

28. (Previously Presented) The method of claim 14, further comprising:

maintaining, in a second data store, a current mirror copy of the digital content of the storage system; and

overwriting the at least one unit of storage with the new data and updating the current mirror copy in the second data store with the new data, wherein the overwriting and the updating occur after the old data is copied to the data store.

29. (Previously Presented) The computer readable medium of claim 23, further comprising:

code adapted to maintain, in a second data store, a current mirror copy of the digital content of the storage system; and

code adapted to overwrite the at least one unit of storage

with the new data and update the current mirror copy in the second data store with the new data after the old data is copied to the data store.